

Neurons in Anterior Cingulate Cortex Multiplex Information about Reward and Action

Benjamin Y. Hayden^{1,2,3} and Michael L. Platt^{1,2,3,4}

¹Department of Neurobiology and ²Center for Neuroeconomic Studies, Duke University School of Medicine, and ³Center for Cognitive Neuroscience and

⁴Department of Evolutionary Anthropology, Duke University, Durham, North Carolina 27708

The dorsal anterior cingulate cortex (dACC) is thought to play a critical role in forming associations between rewards and actions. Currently available physiological data, however, remain inconclusive regarding the question of whether dACC neurons carry information linking particular actions to reward or, instead, encode abstract reward information independent of specific actions. Here we show that firing rates of a majority of dACC neurons in a population studied in an eight-option variably rewarded choice task were sensitive to both saccade direction and reward value. Furthermore, the influences of reward and saccade direction on neuronal activity were approximately equal in magnitude over the range of rewards tested and were statistically independent. Our results indicate that dACC neurons multiplex information about both reward and action, endorsing the idea that this area links motivational outcomes to behavior and undermining the notion that its neurons solely contribute to reward processing in the abstract.

Introduction

A rat in a maze, a hunter–gatherer on the plains, and a football quarterback in the middle of a tense play are each confronted with a choice between several possible actions. Choosing the best action often depends on learning the rewards and punishments associated with each action from experience. Anatomical, lesion, and neurophysiological studies have implicated the anterior cingulate cortex (ACC) in this process. ACC lesions impair outcome–action mapping and reduce behavioral adjustments associated with changing reward contingencies (Hadland et al., 2003; Walton et al., 2003; Rudebeck et al., 2006b, 2008; Kennerley et al., 2006) and can cause deficits that are restricted to specific motor modalities (Turken and Swick, 1999). Firing rates of ACC neurons are selective for specific action–reward combinations (Matsumoto et al., 2003). Moreover, ACC receives projections from structures that process rewards, including the orbitofrontal cortex, the striatum, and the mesolimbic dopamine system, and structures that process actions and locations, including dorsolateral prefrontal cortex (DLPFC) and the supplementary and primary motor cortices (Oades and Halliday, 1987; Vogt and Gabriel, 1993; Gaspar et al., 1995; Paus, 2001).

Many authors have proposed that ACC plays a more general role in outcome monitoring and subsequent behavioral adjustment (Gehring and Willoughby, 2002; Ito et al., 2003; Kerns et al., 2004; Kennerley et al., 2006; Rudebeck et al., 2008). In partial support of this idea, single neurons in dorsal anterior cingulate

cortex (dACC) respond strongly after actions that lead to rewards and signal the size of the obtained rewards (Shidara and Richmond, 2002; Williams et al., 2004; Amiez et al., 2006; Matsumoto et al., 2007; Sallet et al., 2007; Quilodran et al., 2008; Hayden et al., 2009; Kennerley and Wallis, 2009a,b; Luk and Wallis, 2009). Dorsal ACC also appears to signal when obtained rewards fall below a desired level, precipitating a change in action selection (Shima and Tanji, 1998; Williams et al., 2004; Kennerley et al., 2006; Seo and Lee, 2007; Quilodran et al., 2008; Rudebeck et al., 2008; Hayden et al., 2009).

The available data do not discriminate whether ACC represents abstract reward information only or more specific information linking particular actions to outcomes because the selectivity of ACC neurons for specific actions remains almost entirely unknown. To address these questions, we probed the activity of single neurons within dACC while monkeys performed a variably rewarded eight-option decision task. We found that postsaccade responses of most dACC neurons strongly encoded the direction of the saccade. These data suggest that dACC neurons do not solely represent rewards but instead multiplex information about the specific action used to acquire the reward.

Materials and Methods

The data reported in this paper were originally collected for another purpose and published (Hayden et al., 2009). The analyses presented here are all new.

Surgical and behavioral procedures. All procedures were approved by the Duke University Institutional Animal Care and Use Committee and were designed and conducted in compliance with the National Institutes of Health *Guide for the Care and Use of Animals*. Two male rhesus monkeys (*Macaca mulatta*) served as subjects. Both animals were implanted with a small head-holding prosthesis using standard techniques. Six weeks later, they were habituated to laboratory conditions and trained to perform oculomotor tasks for liquid reward. A second procedure was then performed to place a plastic recording chamber over dACC. Animals received analgesics and antibiotics after all surgeries. Throughout

Received Sept. 29, 2009; revised Dec. 16, 2009; accepted Dec. 21, 2009.

This research was supported by National Institutes of Health Grant R01EY013496 and Kirschstein National Research Service Award DA023338 (B.Y.H.). We thank Karli Watson for help in training the animals and Sarah Heilbronner and Steve Chang for useful comments in the design of the experiment and the analysis.

Correspondence should be addressed to Benjamin Y. Hayden, Department of Neurobiology, Duke University Medical School, Durham, NC 27710. E-mail: hayden@neuro.duke.edu.

DOI:10.1523/JNEUROSCI.4874-09.2010

Copyright © 2010 the authors 0270-6474/10/303339-08\$15.00/0

both behavioral and physiological recording sessions, the chamber was kept sterile with regular antibiotic washes and sealed with sterile caps.

Horizontal and vertical eye positions were sampled at 1000 Hz by an infrared eye-monitoring camera system (Eyelink 1000). Stimuli were controlled by a computer running Matlab with the Psychophysics Toolbox (Brainard, 1997) and the Eyelink Toolbox (Cornelissen et al., 2002). Visual stimuli were small colored rectangles ($\sim 1^\circ$) on a computer monitor placed directly in front of the animal and centered on his eyes. A solenoid valve controlled the duration of juice delivery. Reward volume was 0 μl (dark gray square), 100 μl (red square), 200 μl (light gray square), 267 μl (orange square), 300 μl (purple square), 333 μl (blue square), and 367 μl (bright green square).

Every trial began when eight white squares appeared (Fig. 1). The monkey had half a second to inspect these stimuli with unrestrained gaze. After this inspection period, a small yellow square (fixation point) appeared at the center of the screen. Once monkeys had aligned gaze with this square ($\pm 2^\circ$), a half second fixation epoch began. Breaks of fixation (deviations of greater than the 2° fixation window) during this epoch were considered errors and were followed by an immediate timeout and no reward. After the fixation epoch, the central square disappeared and the monkey was free to shift gaze to one of the peripheral targets ($\pm 3^\circ$). Failure to shift gaze to one of these targets within 3 s was considered an error and was followed by an immediate timeout and no reward (although in practice monkeys almost always shifted gaze within a few hundred milliseconds).

Immediately after this saccade, all eight squares changed color. Seven of the targets [the low value (LV) targets] turned red. The other one [the high value (HV) target] turned one of six colors: dark gray, light gray, orange, purple, blue, and bright green (16.67% probability of each). After another half a second, a reward was given corresponding to the color of the selected target (see above). Saccades subsequent to the first one had no effect on reward. The location of the HV target either remained where it was (60% of trials) or moved one square clockwise (40% chance). On all trials, the value of the selected target was considered the “experienced reward.” On trials in which the monkey chose the LV target, we considered the value of the HV target the “fictive reward.” Because neural responses on LV choice trials were influenced by fictive outcomes (Hayden et al., 2009), we avoided this confounding factor by excluding LV choice trials from all analyses in the present study, except when noted.

Microelectrode recording techniques. Single electrodes were lowered with a hydraulic microdrive (David Kopf Instruments) until the waveform of one to three neurons was isolated. Individual action potentials were identified by standard criteria and isolated on a Plexon system. Neurons were selected for recording on the basis of the quality of isolation only and not on task-related response properties. We approached ACC through a standard recording grid. ACC was identified by structural magnetic resonance images taken before the experiment. Neuroimaging was performed at the Center for Advanced Magnetic Development at Duke University Medical Center, on a 3 T MR Imaging Instrument using 1 mm slices. We confirmed that electrodes were in ACC using stereotactic measurements, as well as by listening for characteristic sounds of white and gray matter during recording. Recordings were made between the position of 26 and 30 mm anterior to the interaural plane, with most occurring between 27 and 29. ACC recordings were made in areas 24c (and possibly 6/32) in the cingulate sulcus, corresponding closely to what

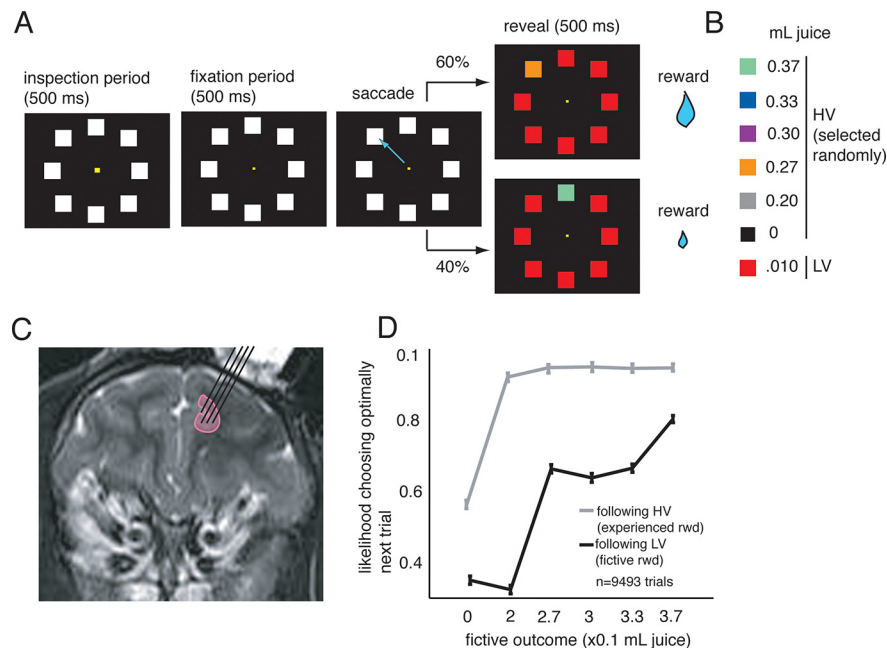


Figure 1. Schematic of eight-target saccade choice task and recording location. **A**, Small yellow fixation point appears, surrounded by eight white squares arranged in a ring. When a monkey fixates on a point, its size shrinks, and after half a second, the monkey is free to shift gaze to one of eight targets. After saccade, all targets change color. Seven (LV targets) turn red, and one (HV target) turns another color. Subsequent saccades have no effect on reward. After a half-second delay, reward is given. **B**, Between trials, HV target either remains at the same position (60% chance) or moves to an adjacent position (40% chance). Value associated with the HV target is indicated by its color. Only HV trials were analyzed in the present study. **C**, Recordings were made in the dorsal and ventral banks of the sulcus of the dACC in two monkeys (magnetic resonance image of monkey E shown). Additional details on recording location can be found in the supplemental data (available at www.jneurosci.org as supplemental material). **D**, Performance of both monkeys on this task. Plot of average likelihood of choosing the optimal target as a function of each possible experienced (gray) and fictive (black) outcomes. For the present study, we focused on HV trials (gray line).

is called dACC. These recordings are somewhat inferior to the location where neurons are found to exhibit reward monitoring signals [sometimes called supplementary motor area (Stuphorn et al., 2000)], superior to an area implicated in social behavior in a lesion study [perigenual, or gyral, ACC (Rudebeck et al., 2006a)], and rostral to the cingulate motor areas (Shima et al., 1991; Shima and Tanji, 1998). Additional details on the recording location can be found in the supplemental data (available at www.jneurosci.org as supplemental material).

Neuronal analyses. Peristimulus time histograms (PSTHs) were constructed by aligning spike rasters to trial events and averaging firing rates across multiple trials. For display, each PSTH was smoothed using a 100 ms running boxcar. Data were aligned to the time of the first saccade made to a target. In one analysis, we compared these responses with those generated in other epochs, as specified in the text. Statistical comparisons were performed on binned, unsmoothed, firing rates of single neurons in 500 ms bins. All analyses were repeated with bins of other sizes (300 ms, 400 ms, 750 ms, and 1 s), and nearly identical results were obtained in all cases.

To quantify spatial tuning functions, we fit cosine curves to plots of firing rate as a function of target position using a simple gradient descent procedure in Matlab. Our fitting program minimized the error between the fitted cosine function and the eight locations using the raw averaged data. Cosine wave frequency was allowed to have one of two values, corresponding to unimodal and bimodal curves. To test for significance of tuning, we used the following bootstrap method. We randomly reshuffled the angle of the saccades and then determined the amplitude of the cosine curve fit to the randomized points using the procedure described above. (For neurons that were found to be bimodal by the fitting procedure, we generated randomized bimodal cosine curves.) We then repeated this process 10,000 times. If the amplitude of the curve fit to the original data was greater than the amplitude of 95% of the randomized curves, we considered the tuning significant at $p < 0.05$; if $>99\%$ of the randomized curves, we considered tuning significant and $p < 0.01$.

In our analysis of the temporal dynamics of directional tuning, we divided trials into eight 500 ms epochs defined relative to the time of the choice. Epoch 1 (pretrial) began 2.5 s before the gaze shift to the fixation spot and ended 500 ms later, well before the appearance of the targets. Epoch 2 (prefixation) began 1.5 s before the choice, ended 500 ms later, and included the presentation of the eight targets. Epoch 3 (fixation) began 1 s before the choice and included fixation on the central spot. Epoch 4 (saccade) began 500 ms before the choice and ended 500 ms later, whereas epoch 5 (postsaccade) began immediately at choice saccade initiation and lasted 500 ms. Epoch 6 (postreward) began 500 ms after choice saccade initiation coincident with the delivery of the juice reward. Epochs 7 and 8 began 1 and 1.5 s after the choice, during the intertrial interval.

For some analyses, we performed a linear regression of firing rate for a given direction against the distance from that direction to the nearest peak of the fit tuning curve. For unimodally tuned neurons, this variable was defined as the distance between target position and the tuned direction of the neuron (the direction that elicited maximal firing), scaled from 0 to 1, with 0 corresponding to 180°. For bimodally tuned neurons, that variable was defined as the distance between target position and the nearest of the two peak positions (scaled from 0 to 1, with 0 corresponding to 90°). If the regression gave a value that was significantly different from 0 (coefficient from regression test), we considered the tuning significant.

Analyses of the influence of reward and saccade direction on firing rates were performed with a simple linear regression. We first regressed firing rates against target position, as above, and reward magnitude. To assess the relative contribution of spatial and reward information, we normalized reward values from 0 (corresponding to 0 μ l) to 1 (corresponding to 333 μ l). In all cases, data from trials with 100 μ l rewards (LV targets) were excluded from analysis. These trials were closely associated with fictive outcomes and may have had particular behavioral relevance for learning (Hayden et al., 2009). However, analyses were repeated with these trials included and treated as 100 μ l trials (i.e., ignoring fictive outcomes), and nearly identical results were obtained in all cases.

For this regression, we used the following equation: $FR = \beta_{sp}s + \beta_{rd}r + \beta_{int}(s \times r)$. We regressed firing rates (FR) against two variables, both ranging from 0 to 1, corresponding to reward size and saccade direction. We scaled reward (r in the equation) from 0 (0 μ l) to 1 (333 μ l) and spatial location (s in the equation) between 1 (at peak position) and 0 (180° from peak position for unimodal neurons and 90° from peak position for bimodal neurons).

Results

Behavior

We recorded responses of 68 neurons in dACC while two monkeys performed an eight-target saccade task. We previously used this task to investigate the role of ACC in fictive learning (Hayden et al., 2009). On each trial, monkeys chose one of eight targets arranged in a circle. One of these targets (the HV target) was associated with a variable reward of higher expected value than the others (the LV target). The computer selected the reward for the HV target randomly from a distribution of six possible rewards (0, 200, 267, 300, 333, and 367 μ l), whereas the rewards for the LV targets were always the same (100 μ l). The position of the HV target was not explicitly cued but depended on a stochastically predictable pattern: it reappeared at the same location with 60% likelihood and one square clockwise with 40% likelihood. Consequently, the optimal strategy on this task was to choose the location that was the HV target on the previous trial (hereafter known as the “optimal target”). Because the HV target location systematically moved clockwise throughout all eight radial positions over the course of each session, we could assess neuronal selectivity to saccade direction.

Monkeys performed the task quite well. They selected the HV target on 48% of trials and chose the optimal target on 75% of trials (Fig. 1*D*). Deviations from the optimal target were more often adjacent to it than farther away (binomial test, $p < 0.001$),

further confirming that monkeys generally understood the task. For the purposes of the present study, monkeys were equally likely to choose each of the eight targets during each recording session, because each of the eight targets was equally likely to serve as the HV target across a session (χ^2 test, $p = 0.7$). Moreover, the average reward obtained from each target was identical (one-way ANOVA, average $p = 0.25$). Reaction times did not depend on saccade direction (ANOVA, average $p = 0.44$) or reward size (ANOVA, average $p = 0.35$). These results suggest that differential firing rates when saccades were directed to different target locations (see below) did not reflect intervening variables such as arousal or attention.

ACC neurons are selective for saccade direction

Neurons in dACC respond strongly and phasically around the time of saccades to rewarding targets (Amiez et al., 2006; Matsumoto et al., 2007; Sallet et al., 2007; Hayden et al., 2009; Kennerley et al., 2009). We therefore first focused on these phasic responses in our analysis of saccade tuning. In this task, impending reward size was revealed as soon as the monkey aligned gaze with a target, although reward was not delivered for another 500 ms. We analyzed firing rates of all neurons during the 500 ms epoch beginning after the reward amount was revealed and ending before the reward was obtained (postsaccade epoch).

For example, responses of one neuron were greatest around the time of contraversive saccades but were less enhanced around the time of ipsiversive saccades (Fig. 2*A*). We analyzed responses of this neuron during the epoch indicated in gray. Figure 2*B* (solid line) shows the average firing rate of a single dACC neuron after gaze shifts to each of the eight targets used in our task. The same raw data are plotted again in a polar format in Figure 2*B*. This neuron was sensitive to saccade direction, with greatest firing rates after saccades in the 45° (upward, contraversive) direction and lowest firing rates after saccades in the opposite (225°) direction.

We next calculated a best-fit cosine function for the responses of this neuron as a function of saccade direction (Fig. 2*B*, dashed line). The mean firing rate for this neuron was 38.1 spikes/s, whereas the fit amplitude (peak minus trough of curve) was 28 spikes/s. The peak of the tuning curve was found to be 38°, a close match to the target location with highest firing rate (i.e., 45°). The cosine curve provided a significantly better fit than a horizontal line drawn through the mean firing rate of the neuron (bootstrap test, $p < 0.01$; see Materials and Methods). The bootstrapping procedure is designed to provide a measure of significance, whereas the firing rate to amplitude ratio provides one measure of effect size.

Figure 2*D* shows the average firing rate for a different single neuron after saccades to each of the eight targets used in our task (solid, raw data; dashed, fit data). The same raw data are plotted in polar format in Figure 2*E*. This neuron showed bimodal tuning for target position, exhibiting greatest firing rates after gaze shifts to targets at 90 and 315° and lowest firing rates after gaze shifts to targets at 0 and 180°. We then calculated the best-fit single-amplitude single cosine curve to these data. The mean firing rate for this neuron was 31.2 spikes/s, whereas the fit amplitude (peak minus trough) was 32.3 spikes/s and the fitted phase was 94°. Our algorithm fit this neuron with a bimodal tuning curve. The cosine curve provided a significantly better fit than a horizontal line drawn through the mean firing rate of the neuron (bootstrap test, $p < 0.01$).

We observed significant saccade direction tuning in 56 of 68 neurons (82.3%), whereas 3.4 (5%) neurons would be expected

to show such tuning by chance. The effects we observed were consistent in both monkeys. In monkey E, we observed significant tuning in 43 of 50 (86%) neurons, whereas 2.5 (5%) neurons would be expected to show such tuning by chance. In monkey O, we observed significant tuning in 13 of 18 (72.2%) neurons, whereas 0.9 (5%) neurons would be expected to show such tuning by chance.

The majority of significantly tuned neurons ($n = 38$, 68%) showed unimodal tuning, whereas a minority ($n = 18$, 32%) showed bimodal tuning for saccade direction. The average firing rate of all neurons during the postsaccade epoch was 19.22, whereas the average amplitude of these neurons was 21.76, meaning that the average ratio of amplitude to mean was 1.13 (Fig. 3A). A plot of the peak tuning direction for the unimodal neurons is shown in Figure 3B. This distribution is significantly biased toward contraversive saccades (binomial test, $p < 0.001$). We note that this contraversive bias is mediated primarily by a preponderance of neurons tuned for 45° . This effect may reflect sampling bias within the subregion of dACC in which we recorded; future studies will be needed to delineate the topography of spatial tuning in dACC. Also, we note that neurons may equally be said to be selective for target position and saccade direction; our data do not distinguish these possibilities.

Temporal dynamics of saccade direction and reward on firing rate

Neurons in this task exhibited their strongest phasic responses around the time of the gaze shift to a target for reward (Fig. 4A). We next investigated the temporal dynamics of saccade direction tuning during the course of the task. Figure 4B shows the saccade direction tuning curve of a single neuron (same as in Fig. 2A,B) during several 500 ms task epochs. This neuron showed no saccade direction tuning before the trial (epoch 1, bootstrap test, $p = 0.19$), suggesting no impending directional bias (cf. Lauwereyns et al., 2002). Weak directional tuning emerged around the time of fixation (epochs 2 and 3) and peaked around the time of the saccade (epochs 4 and 5). Note that, although reward information was not specified during epoch 4 and fully specified during epoch 5, saccade direction tuning did not differ during these two epochs (bootstrap test, $p = 0.14$), consistent with the notion that reward tuning and location tuning do not strongly interact in this task (see below). After reward delivery, direction tun-

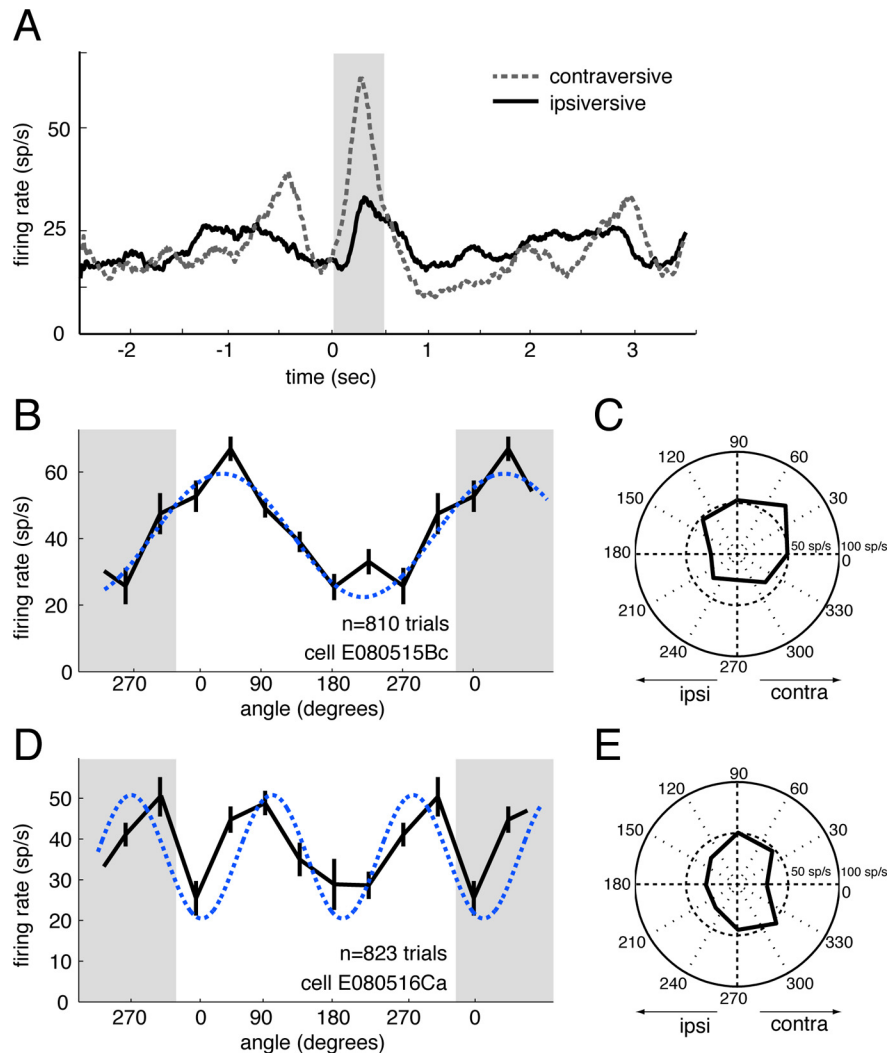


Figure 2. Tuning curves for two single neurons from our sample. **A**, Plot of firing rate of a single neuron on trials with contraversive saccades (dashed gray) and ipsiversive saccades (black). Data are aligned to the onset of the saccade (time 0 on the graph). **B**, Average firing rate of example dACC neuron during 300 ms epoch after saccades to rewarding targets depends on target position (black lines). Bars indicate 1 SE. Dashed line indicates values for best-fit sine-wave curve. Gray areas represent points that are repeated to emphasize the sinusoidal pattern. **C**, Firing rate data replotted in polar coordinate system. This neuron fired maximally after saccades to the top right quadrant (contraversive saccades). **D**, Average firing rate of a different neuron that showed bimodal tuning. **E**, Firing rate shown in polar coordinate system.

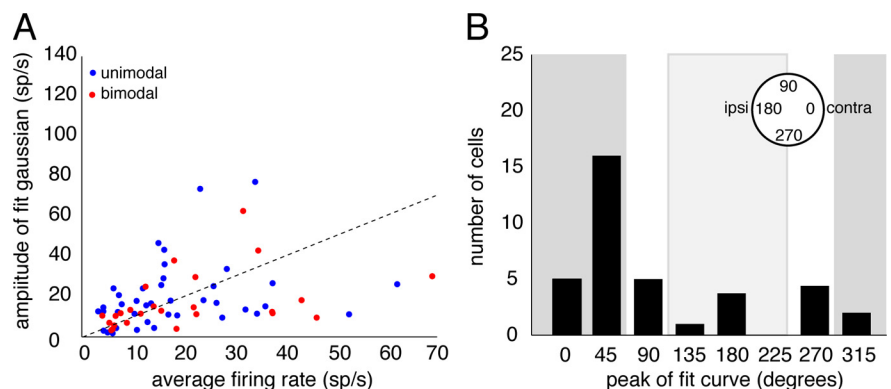


Figure 3. Saccade direction tuning in the dACC neuronal population. **A**, Scatter plot of average firing rate of all neurons in sample versus amplitude of sinusoid fit to their firing rates. Dashed line indicates unity. **B**, Bar plot of cosine function peak fitted to neuronal tuning functions. Black, Unimodal neurons; white, bimodal neurons. For bimodal neurons, the larger peak is shown. Neurons are more likely to be tuned for contraversive actions (dark gray boxes) than ipsiversive actions (light gray box).

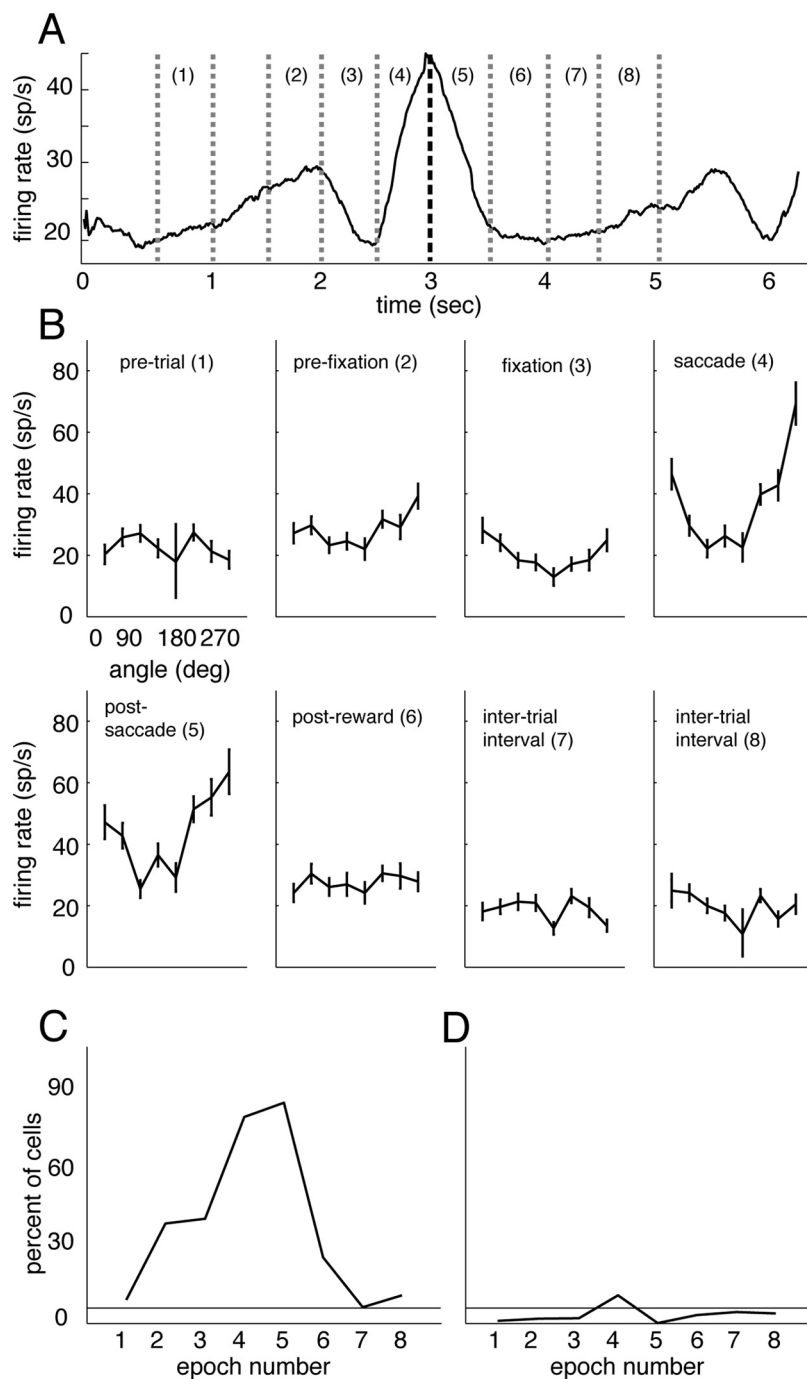


Figure 4. Temporal dynamics of saccade direction tuning in dACC neurons. **A**, Plot of average firing rate of single dACC neuron across critical points in task. Same neuron as in Figure 2. Dashed lines indicate epochs used in analysis. Saccade to target occurs between epochs 4 and 5 (thick dashed line). **B**, Plot of direction tuning curves for this neuron during each of the epochs analyzed. Bars indicate 1 SE in firing rate. Neuron was weakly tuned for direction during early task epochs, strongly tuned during perisaccadic epochs, and not tuned for saccade direction before and after trials. **C**, Plot of percentage of neurons showing significant tuning in different task epochs. Horizontal line, Proportion of neurons expected by chance. **D**, Proportion of neurons exhibiting significant spatial tuning in a direction that differed from the tuning direction observed in the postsaccadic epoch (epoch 5).

ing disappeared (epoch 6) and remained absent during the early and late phases of the intertrial interval (epochs 7 and 8).

We observed similar patterns of spatial tuning dynamics at the population level. The number of neurons exhibiting significant saccade direction tuning effect is shown in Figure 4C. These data demonstrate that the population was most strongly selective for saccade direction immediately before and after saccade initiation

and not at other times. Our data also indicate that saccade direction tuning does not generally vary across epochs: Figure 4D shows a plot of the number of neurons in each epoch that were significantly tuned for a saccade direction different from the saccade direction tuning observed in the postsaccadic epoch (i.e., epoch 5). These data demonstrate that the number of neurons exhibiting a significant change in tuning direction is below that expected by chance ($n = 3.4$ of 68, 5%) in all epochs except epoch 4 and only slightly above chance in that epoch ($n = 6$ of 68, 8.8%).

Interaction of saccade direction tuning with reward modulation

We next compared coding of saccade direction in dACC with coding of reward size. We previously reported that 46 of 68 (67%) neurons in this dataset were selective for reward size during a 300 ms epoch beginning 100 ms after the end of the saccade, although we did not previously examine spatial tuning (Hayden et al., 2009). Here we report that 40 (59%) of the neurons in our population were tuned for both saccade direction and reward size, whereas six were tuned for neither variable.

Figure 5A shows the directional tuning curves of a single neuron separately for large (300 and 333 μ l), small (200 μ l), and 0-size (0 μ l) rewards. These data suggest that the main effects of reward and saccade direction for this neuron do not interact and are thus statistically independent (two-way ANOVA, firing rate against reward size and target direction, $p = 0.5$ for the interaction of the two main effects). To examine this issue more closely, we compared effect sizes in the two conditions (Fig. 5B). We defined the effect size of rewards as the difference in firing for the large and small rewards; we defined the effect size of direction as the difference in firing for direction closest to the fit peak of the tuning curve and the direction 180° away (90° for the bimodal neurons).

To facilitate comparison between these variables at the population level, we calculated regression coefficients expressing the dependence of firing rates on reward size and target position separately. We

also calculated an interaction term. The mean regression coefficient for reward size was 0.27, whereas the average regression coefficient for target location was 0.19. The median coefficient for reward size was 0.14, whereas the median for target location was 0.15 (these are not significantly different, Wilcoxon's rank sum test, $p = 0.2$). We conclude that the extent of encoding for saccade direction is approximately as strong as that for reward size,

at least within the range of rewards and saccade directions used in this task.

To more closely examine whether neurons that were well tuned for reward were less well tuned for direction, we calculated a correlation coefficient for the population dataset of reward and direction correlation coefficients for each neuron. A negative correlation coefficient would indicate that tuning strength for reward value is anticorrelated with tuning for saccade direction. We obtained a value of $r^2 = -0.00000036$, a value that is not significant ($p > 0.5$). Among neurons significantly modulated by both value and direction, the value of the correlation coefficient was $r^2 = 0.00002$, a value that is also not significant ($p > 0.5$).

Of the 68 neurons in our population, we observed a statistically significant interaction between reward size and saccade direction in 11 (16.9%, ANOVA, $p < 0.05$ for interaction term). For these 11 neurons, the size of the regression coefficient for the interaction was much smaller than the size of the main effects (mean coefficients for these neurons were 3.9 for the interaction, 31.2 for position, and 25.0 for reward). These data endorse the notion that interaction between saccade direction and reward size is relatively uncommon in ACC and that reward and location are generally multiplexed in independent, additive channels.

On approximately half the trials in this task, monkeys chose an LV target and learned the value of the outcome they would have obtained had they chosen the HV target. We previously showed that ACC neurons respond to fictive outcomes as well as experienced outcomes and generally show monotonically increasing firing rates for larger fictive outcomes (Hayden et al., 2009). Responses on these LV trials were not included in the previous analyses. We next examined whether coding for fictive outcomes interacts statistically with coding for target location. To examine this question, we performed a multiple linear regression of firing rates of ACC neurons on fictive reward size and saccade target location as defined above. Of the 68 neurons in our population, we observed a statistically significant interaction in four (5.8%, ANOVA, $p < 0.05$ for interaction), a proportion that is similar to that expected by chance (i.e., 5%). For these four neurons, the size of the regression coefficient for the interaction was much smaller than the size of the main effects (mean coefficients for these neurons were 4.8 for the interaction, 20.9 for position, and 14.1 for fictive reward).

Discussion

Here we report that firing rates of dACC neurons depend strongly on saccade direction in an eight-option choice task. Saccade direction tuning was observed weakly during decision and motor planning stages, strongly before and after the gaze shift to the location, and then rapidly disappeared within 500 ms after the reward was given. The relative contribution of target location and reward information to neuronal firing rates was approximately the same, and the two effects interacted only weakly.

Studies of ACC function have tended to emphasize either the abstract (i.e., effector independent) nature of reward encoding (Procyk et al., 2000; Williams et al., 2004; Amiez et al., 2005, 2006;

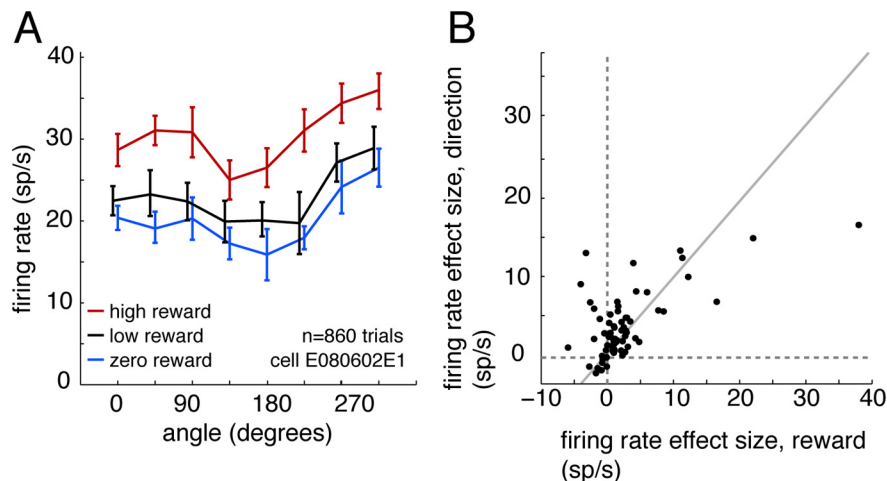


Figure 5. Saccade direction and reward size are encoded independently by dACC neurons. **A**, Plot of average firing rate of an example neuron in zero (blue line), low (black line), and high (red line) reward conditions. Although neuron represents both saccade direction and reward size, these two variables are represented independently (i.e., reward does not modulate the gain of the saccade response). **B**, Plot of effect size for reward (abscissa, response to largest minus response to smallest reward) versus effect size for saccade direction (ordinate, response to peak direction minus response to trough direction). Diagonal line indicates unity.

Matsumoto et al., 2007; Sallet et al., 2007; Seo and Lee, 2007; Quilodran et al., 2008; Hayden et al., 2009; Kennerley et al., 2009) or the role of this area in associating specific locations and actions with rewards (Turken and Swick, 1999; Hadland et al., 2003; Rushworth et al., 2004; Rudebeck et al., 2008). Our results support both points of view. The co-occurrence of tuning for saccade direction/target location and reward size suggests a potential mechanism by which abstract reward information may be integrated into long-term action–reward associations. We conjecture that, during action–outcome learning, neuronal activity associated with greater rewards strengthens connections between action-specific neurons and their associated actions. Future studies will be needed to confirm this hypothesis.

Some previous studies have shown weak or inconsistent tuning for action in dACC (Ito et al., 2003; Matsumoto et al., 2007; Seo and Lee, 2007). Other studies have reported clear, but less common, tuning for actions (Matsumoto et al., 2003; Kennerley and Wallis, 2009; Luk and Wallis, 2009). We hypothesize that the dynamic nature of our task, which incentivized learning from experienced and fictive reward outcomes, may have enhanced representation of information about the specifics of actions leading to reward (for similar ideas, see Luk and Wallis, 2009). Notably, our finding of value-independent response-selective tuning around the time of the reward is reminiscent of a recent study showing similar effects in the orbitofrontal cortex, suggesting that response tuning may be more common than previously thought (Tsujimoto et al., 2009).

The importance of saccade direction in shaping the firing rates of neurons in dACC aligns this area with others that show similar action tuning. For example, neurons in parietal cortex are tuned for attentional locus, direction of planned saccades, and strength of evidence in favor of a specific saccade (Andersen and Buneo, 2002; Gold and Shadlen, 2002; Bisley and Goldberg, 2003); DLPFC and the mediodorsal nucleus of the thalamus are tuned for remembered locations (Funahashi et al., 1989; Watanabe and Funahashi, 2004), and the frontal and supplementary eye fields are tuned for saccade direction (Schall, 1991; Schall and Hanes, 1993). It is difficult for us to compare the size of saccade direction tuning in ACC with that observed in other areas because previous

studies typically preselected neurons on the basis of task response properties. Nonetheless, our data suggest that the extent of spatial tuning within the dACC is not substantially different from other brain regions outside the visual and motor cortices, including lateral intraparietal cortex, the frontal eye fields, and superior colliculus (Schall, 1991; Platt and Glimcher, 1998; Andersen and Buneo, 2002; Gold and Shadlen, 2002; Bisley and Goldberg, 2003). In any case, spatial tuning appears to be stronger and more robust than in a main target area of the dACC, posterior cingulate cortex (Olson et al., 1996; Dean et al., 2004; Dean and Platt, 2006).

The ACC is a functionally heterogeneous region, variously divided into several subregions (Vogt and Gabriel, 1993; Paus et al., 1996; Paus, 2001; Lee et al., 2007). Our data were collected only from the cortex surrounding the sulcus of the anterior cingulum (especially the dorsal bank of the sulcus), rostral to the cingulate motor areas, and superior to the perigenual and subgenual cingulate cortices, including the gyrus of the ACC (supplemental Fig. S2, available at www.jneurosci.org as supplemental material). This area is close to the same sites that have been recorded from, and lesioned, in many previous studies that established this area as a critical nexus for reward-based decision making (Procyk et al., 2000; Amiez et al., 2005, 2006; Kennerley et al., 2006, 2009; Sallet et al., 2007; Seo and Lee, 2007; Quilodran et al., 2008; Rudebeck et al., 2008). Given the strong connections between this region and the DLPFC (Vogt et al., 1992; Vogt and Gabriel, 1993), dACC neurons may be more strongly spatially tuned than neurons in more ventral aspects of the frontal lobe. Thus, quantifying the spatial selectivity of neurons throughout the ACC remains an important future task for understanding how these areas contribute to learning, planning, and decision making.

References

- Amiez C, Joseph JP, Procyk E (2005) Anterior cingulate error-related activity is modulated by predicted reward. *Eur J Neurosci* 21:3447–3452.
- Amiez C, Joseph JP, Procyk E (2006) Reward encoding in the monkey anterior cingulate cortex. *Cereb Cortex* 16:1040–1055.
- Andersen RA, Buneo CA (2002) Intentional maps in posterior parietal cortex. *Annu Rev Neurosci* 25:189–220.
- Bisley JW, Goldberg ME (2003) Neuronal activity in the lateral intraparietal area and spatial attention. *Science* 299:81–86.
- Brainard DH (1997) The psychophysics toolbox. *Spat Vis* 10:433–436.
- Cornelissen FW, Peters EM, Palmer J (2002) The Eyelink Toolbox: eye tracking with MATLAB and the Psychophysics Toolbox. *Behav Res Methods Instrum Comput* 34:613–617.
- Dean HL, Platt ML (2006) Allocentric spatial referencing of neuronal activity in macaque posterior cingulate cortex. *J Neurosci* 26:1117–1127.
- Dean HL, Crowley JC, Platt ML (2004) Visual and saccade-related activity in macaque posterior cingulate cortex. *J Neurophysiol* 92:3056–3068.
- Funahashi S, Bruce CJ, Goldman-Rakic PS (1989) Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J Neurophysiol* 61:331–349.
- Gaspar P, Bloch B, Le Moine C (1995) D1 and D2 receptor gene expression in the rat frontal cortex: cellular localization in different classes of efferent neurons. *Eur J Neurosci* 7:1050–1063.
- Gehring WJ, Willoughby AR (2002) The medial frontal cortex and the rapid processing of monetary gains and losses. *Science* 295:2279–2282.
- Gold JI, Shadlen MN (2002) Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* 36:299–308.
- Hadland KA, Rushworth MF, Gaffan D, Passingham RE (2003) The anterior cingulate and reward-guided selection of actions. *J Neurophysiol* 89:1161–1164.
- Hayden BY, Pearson JM, Platt ML (2009) Fictive reward signals in the anterior cingulate cortex. *Science* 324:948–950.
- Ito S, Stuphorn V, Brown JW, Schall JD (2003) Performance monitoring by the anterior cingulate cortex during saccade countermanding. *Science* 302:120–122.
- Kennerley SW, Wallis JD (2009a) Evaluating choices by single neurons in the frontal lobe: outcome value encoded across multiple decision variables. *Eur J Neurosci* 29:2061–2073.
- Kennerley SW, Wallis JD (2009b) Encoding of reward and space during a working memory task in the orbitofrontal cortex and anterior cingulate sulcus. *J Neurophysiol* 102:3352–3364.
- Kennerley SW, Walton ME, Behrens TE, Buckley MJ, Rushworth MF (2006) Optimal decision making and the anterior cingulate cortex. *Nat Neurosci* 9:940–947.
- Kennerley SW, Dahmubed AF, Lara AH, Wallis JD (2009) Neurons in the frontal lobe encode the value of multiple decision variables. *J Cogn Neurosci* 21:1162–1178.
- Kerns JG, Cohen JD, MacDonald AW 3rd, Cho RY, Stenger VA, Carter CS (2004) Anterior cingulate conflict monitoring and adjustments in control. *Science* 303:1023–1026.
- Lauwereyns J, Watanabe K, Coe B, Hikosaka O (2002) A neural correlate of response bias in monkey caudate nucleus. *Nature* 418:413–417.
- Lee D, Rushworth MF, Walton ME, Watanabe M, Sakagami M (2007) Functional specialization of the primate frontal cortex during decision making. *J Neurosci* 27:8170–8173.
- Luk CH, Wallis JD (2009) Dynamic encoding of responses and outcomes by neurons in medial prefrontal cortex. *J Neurosci* 29:7526–7539.
- Matsumoto K, Suzuki W, Tanaka K (2003) Neuronal correlates of goal-based motor selection in the prefrontal cortex. *Science* 301:229–232.
- Matsumoto M, Matsumoto K, Abe H, Tanaka K (2007) Medial prefrontal cell activity signaling prediction errors of action values. *Nat Neurosci* 10:647–656.
- Oades RD, Halliday GM (1987) Ventral tegmental (A10) system: neurobiology. 1. Anatomy and connectivity. *Brain Res* 434:117–165.
- Olson CR, Musil SY, Goldberg ME (1996) Single neurons in posterior cingulate cortex of behaving macaque: eye movement signals. *J Neurophysiol* 76:3285–3300.
- Paus T (2001) Primate anterior cingulate cortex: where motor control, drive and cognition interface. *Nat Rev Neurosci* 2:417–424.
- Paus T, Tomaiuolo F, Otaky N, MacDonald D, Petrides M, Atlas J, Morris R, Evans AC (1996) Human cingulate and paracingulate sulci: pattern, variability, asymmetry, and probabilistic map. *Cereb Cortex* 6:207–214.
- Platt ML, Glimcher PW (1998) Response fields of intraparietal neurons quantified with multiple saccadic targets. *Exp Brain Res* 121:65–75.
- Procyk E, Tanaka YL, Joseph JP (2000) Anterior cingulate activity during routine and non-routine sequential behaviors in macaques. *Nat Neurosci* 3:502–508.
- Quilodran R, Roth é M, Procyk E (2008) Behavioral shifts and action valuation in the anterior cingulate cortex. *Neuron* 57:314–325.
- Rudebeck PH, Buckley MJ, Walton ME, Rushworth MF (2006a) A role for the macaque anterior cingulate gyrus in social valuation. *Science* 313:1310–1312.
- Rudebeck PH, Walton ME, Smyth AN, Bannerman DM, Rushworth MF (2006b) Separate neural pathways process different decision costs. *Nat Neurosci* 9:1161–1168.
- Rudebeck PH, Behrens TE, Kennerley SW, Baxter MG, Buckley MJ, Walton ME, Rushworth MF (2008) Frontal cortex subregions play distinct roles in choices between actions and stimuli. *J Neurosci* 28:13775–13785.
- Rushworth MF, Walton ME, Kennerley SW, Bannerman DM (2004) Action sets and decisions in the medial frontal cortex. *Trends Cogn Sci* 8:410–417.
- Sallet J, Quilodran R, Roth é M, Vezoli J, Joseph JP, Procyk E (2007) Expectations, gains, and losses in the anterior cingulate cortex. *Cogn Affect Behav Neurosci* 7:327–336.
- Schall JD (1991) Neuronal activity related to visually guided saccadic eye movements in the supplementary motor area of rhesus monkeys. *J Neurophysiol* 66:530–558.
- Schall JD, Hanes DP (1993) Neural basis of saccade target selection in frontal eye field during visual search. *Nature* 366:467–469.
- Seo H, Lee D (2007) Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. *J Neurosci* 27:8366–8377.
- Shidara M, Richmond BJ (2002) Anterior cingulate: single neuronal signals related to degree of reward expectancy. *Science* 296:1709–1711.

- Shima K, Tanji J (1998) Role for cingulate motor area cells in voluntary movement selection based on reward. *Science* 282:1335–1338.
- Shima K, Aya K, Mushiake H, Inase M, Aizawa H, Tanji J (1991) Two movement-related foci in the primate cingulate cortex observed in signal-triggered and self-paced forelimb movements. *J Neurophysiol* 65:188–202.
- Stuphorn V, Taylor TL, Schall JD (2000) Performance monitoring by the supplementary eye field. *Nature* 408:857–860.
- Tsujimoto S, Genovesio A, Wise SP (2009) Monkey orbitofrontal cortex encodes response choices near feedback time. *J Neurosci* 29:2569–2574.
- Turken AU, Swick D (1999) Response selection in the human anterior cingulate cortex. *Nat Neurosci* 2:920–924.
- Vogt BA, Gabriel M (1993) *Neurobiology of cingulate cortex and limbic thalamus*. Boston: Birkhauser.
- Vogt BA, Finch DM, Olson CR (1992) Functional heterogeneity in cingulate cortex: the anterior executive and posterior evaluative regions. *Cereb Cortex* 2:435–443.
- Walton ME, Bannerman DM, Alterescu K, Rushworth MF (2003) Functional specialization within medial frontal cortex of the anterior cingulate for evaluating effort-related decisions. *J Neurosci* 23:6475–6479.
- Watanabe Y, Funahashi S (2004) Neuronal activity throughout the primate mediodorsal nucleus of the thalamus during oculomotor delayed-responses. I. Cue-, delay-, and response-period activity. *J Neurophysiol* 92:1738–1755.
- Williams ZM, Bush G, Rauch SL, Cosgrove GR, Eskandar EN (2004) Human anterior cingulate neurons and the integration of monetary reward with motor responses. *Nat Neurosci* 7:1370–1375.